

Abstract of the Disclosure

The present invention includes a scheduling mechanism that fairly allocates a resource to a number of schedulable elements of which some are latency-sensitive. The invention tracks each element's use of the resource by determining the element's virtual time. An active element is selected from the elements that are ready to use the resource by determining the element that has the smallest effective virtual time. The effective virtual time is the element's actual virtual time modified by a borrowed virtual time value. When an element has a short-term need for the resource, it can borrow the privilege to run by borrowing virtual time. As the element uses the resource, it consumes virtual time according to its weight.

When the elements are scheduled for the resource, the ready element having the smallest virtual time is selected. The invention enforces long-term fairness to each element while allowing latency-sensitive elements to be preferably selected. One preferred embodiment is a thread-of-execution scheduler for a computer system. This embodiment allocates the CPU time (the resource) between multiple threads (the elements). Latency-sensitive threads (for example, video display programs) can borrow virtual time so that they are more likely to be scheduled when a condition is satisfied (for example, expiration of a timer that triggers the display of the next video frame). Another preferred embodiment is a queue scheduler for an output-queued data switch. Here multiple queues feed an output port. The multiple queues have different service requirements. The invention provides long-term fairness while still satisfying queues that contain latency-sensitive data.